APPLICATION FOR UNITED STATES PATENT

For

Method and Apparatus For Egress Channel Architecture That Supports Protection Within SONET/SDH Based Networks

Inventors:

Jim Mao

Wei Wu

Prepared by:

BLAKELY SOKOLOFF TAYLOR & ZAFMAN LLP
12400 Wilshire Boulevard
Los Angeles, CA 90025-1026
(408) 720-8300

Attorney's Docket No.: 05043P024

"Express Mail" mailing label number: EL617178485US
Date of Deposit: June 4, 2001
I hereby certify that I am causing this paper or fee to be deposited with the United States Postal Service "Express Mail Post Office to Addressee" service on the date indicated above and that this paper or fee has been addressed to the Commissioner for Patents,
Washington, D. C. 20231

Leah Resendez
(Typed or printed name of person mailing paper or fee)

Leah Resendez
(Signature of person mailing paper or fee)

6-4-01
(Date signed)

METHOD AND APPARATUS FOR EGRESS CHANNEL ARCHITECTURE THAT SUPPORTS PROTECTION WITHIN SONET/SDH BASED NETWORKS

FIELD OF THE INVENTION

[0001] The field of invention relates generally to communication; and more specifically, to a method and apparatus for an egress channel architecture that supports protection within SONET/SDH based networks.

BACKGROUND

<u>Protection within SONET Networks</u>

[0002] Synchronous Optical NETwork (SONET) and Synchronous Digital Hierarchy (SDH) based networks typically emphasize redundancy. That is for example, should a particular network line that couples a pair of networking systems (which may also be referred to as "nodes") within the network fail (or degrade), the network is designed to "switch over" to another network line so that traffic flow is not substantially interrupted. Various types of redundancy may be designed into a SONET network. Some examples are illustrated in the discussion that follows.

[0003] Figure 1 shows a point-to-point perspective. Point to point redundancy focuses on the behavior of a pair of nodes 131, 132 that are coupled together by a plurality of SONET lines 104_1, 104_2, ...104_{x-1}, 104_x. Although other point-to-point schemes may be possible, common point-to-point schemes typically include 1+1 and 1:N. Both schemes classify a network line as either a working

line or a protection line. A working line is deemed as the "active" line that carries the information transported by the network. A protection line serves as a "back-up" for a working line. That is, if a working line fails (or degrades), the protection line is used to carry the working line's traffic.

[0004] In a 1+1 scheme, both the working and protection lines simultaneously carry the same traffic. For example, referring to Figure 1, if line 104_1 is the working line and line 104_2 is the protection line; the transmitting node 131 simultaneously transmits the same information on both the working line 104_1 and the protection line 104_2. The receiving node 132, during normal operation, "looks to" the working line 104_1 for incoming traffic and ignores the incoming traffic on the protection line 104_2. If a failure or degradation of the working line 104_1 is detected, the receiving node 132 simply "looks to" the protection line 104_2 for the incoming traffic (rather than the working line 104_1).

[0005] In a 1:N scheme one protection line backs up N working lines (where N is an integer greater than or equal to 1). For example, referring to Figure 1, lines 104_1 through 104_{x-1} may be established as the working lines while line 104_x may be established as the protection line. If any of the working lines 104_1 through 104_{x-1} fail or degrade, the transmitting node 131 sends the traffic of the failed/degraded working line over the protection line 104_x. The receiving node 132 also "looks to" the protection line 104_x for the traffic that would have been sent over the failed/degraded working line prior to its failure/degradation.

[0006] Figure 2 shows a ring perspective. Ring redundancy schemes focus on the behavior of a plurality of nodes 231 through 234 coupled together in a

ring. Redundancy is commonly handled by sending identical streams of traffic in opposite directions. A first direction may be referred to as the working direction while a second direction may be referred to as the protection direction. The most commonly used types of ring protection are Unidirectional Path Switched Ring (UPSR) and Bi-directional Line Switched Ring (BLSR). In a two-fiber UPSR approach, working traffic is sent in a first direction around the ring (e.g., clockwise) and protection traffic is sent in a second direction around the ring (e.g., counter - clockwise).

[0007] In a Bi-directional Line Switched Ring (BLSR) approach, each pair of rings are viewed as comprising an “inner” ring and “outer” ring (although note that the rings’ actual geographic coverage does not necessarily have to correspond to the inner ring always being within the outer ring). Typically, for each ring, half of the capacity is allocated for working traffic and the other half of the capacity is allocated for protection traffic. As such, both working traffic and protection traffic flow bi-directionally. In either the UPSR or BLSR approaches, if failure or degradation occurs in the working direction, active traffic is looked for in the protection direction.

[0008] More sophisticated SONET networks may also be designed that provide protection at higher degrees of resolution. That is, each SONET line (such as line 104_1 of Figure 1 or line 204 of Figure 2) may be viewed as transporting a number of STS-1 signals. For example, if lines 104_1 and 204 each correspond to an STS-n line, each of these lines may be viewed as carrying n STS-1 signals (e.g., an STS-48 line may be viewed as carrying 48 STS-1

signals). Furthermore, in other environments, each STS-1 signal is used as a resource for carrying a plurality of lower speed signals.

[0009] Protection may be provided for STS-1signals individually or for their constituent lower speed signals individually. Either of these forms of protection are commonly referred to as “path protection”. For example, in one type of 1+1 path protection scheme, an individual “working” STS-1 signal within an STS-n line (rather than all the STS-1 signals on the STS-n line) is backed up by a “protection” STS-1 signal transported on another STS-n line.

[0010] Automatic Protection Switching (APS) is a protocol that may be executed by the nodes within a SONET network. APS allows SONET nodes to communicate and organize the switching over from their working configuration to a protection configuration in light of a failure or degradation event (and then back again after the failure/degradation is corrected). For example, in a typical approach, K1 and K2 bytes are embedded within the SONET frame that is communicated between a pair of nodes in order to communicate failure/degradation events, requests for a switch over, correction thereafter, etc.

<u>Distributed Switch Architecture</u>

[0011] Figure 3 shows an embodiment of a distributed “full mesh” node (or system) architecture 331. The architecture 331 of Figure 3 may be utilized to implement a SONET node such as nodes 131, 132 of Figure 1 or nodes 231 through 234 of Figure 2. An ingress channel receives incoming data from a networking line. Figure 3 shows ingress channels 301_1 through 301_x that each receive incoming data on a respective network line 303_1 through 303_x.

[0012] An egress channel transmits outgoing data onto a networking line. Figure 3 shows egress channels 312_1 through 312_x that each transmit outgoing data on a respective network line 304_1 through 304_x. In a full mesh architecture embodiment, each ingress channel 301_1 through 301_x transmits all of its ingress traffic to each egress channel 312_1 through 312_x. For example, referring to Figure 3, ingress channel 301_1 receives n STS-1 signals from its corresponding network line 303_1 (e.g., if network line 303_1 is an OC-48 line; n=48 and the ingress line channel receives 48 STS-1 signals).

[0013] All n of the STS-1 signals received by the ingress channel 301_1 are transmitted across the node's backplane 305 over each of its output lines 306, 310, 311, 312. A backplane is a board (e.g., a PC or "planar" board) having signal lines that electrically couple various line cards together. Typically, individual cards "plug into" a backplane 305 (e.g., via a card connector) and; in so doing, become communicatively coupled with one another. As such, plugging a sufficient type and quantity of cards into a backplane results in the formation of a networking system. Frequently (although not a strict requirement), backplanes have little or no sophisticated circuitry (e.g., processors, Application Specific Integrated Circuits (ASICs), etc.) and are substantially a collection of "short circuits" from card to card (although passive devices (e.g., capacitors, resistors, etc.), line drivers and other signal enhancing chips or devices are often found on a typical backplane).

[0014] Continuing with the discussion of the full mesh embodiment of Figure 3, note that each egress channel 312_1 through 312_x receives all n STS-1 signals

received by ingress channel 303_1. In one approach, each STS-1 signal is provided its own signal line to each egress channel. As a result, each output 306, 310, 311, 312 corresponds to a n-wide bus. As each ingress channel is similarly designed, each egress channel 312_1 through 312_x receives all the incoming traffic received by the node. For example, in the particular example of Figure 3, there are x ingress channels 301_1 through 301_x that each receive n STS-1 signal. As such, each egress channel 312_1 through 312_x receives xn STS-1 signals (which correspond to the total amount of traffic received by the node 331).

[0015] For example, note that egress channel 312_1 receives inputs 306_1 through 306_x where each of these inputs correspond to the n STS-1 signals received by their corresponding ingress channel (i.e., input 306 for ingress channel 301_1, input 307 for ingress channel 301_2, input 308 for ingress channel 301_3,... and input 309 for ingress channel 301_x). In order to implement the switching fabric of the node, each egress channel 312_1 through 312_x is configured to select n of the xn STS-1 signals and transmit the n STS-1 signals over its corresponding outgoing networking line 304_1 through 304_x.

[0016] Note that there is a distinction between a line card and a channel. A line card is a card that can be coupled to one or more network lines. A channel is a data path within a line card that handles traffic flow in a particular direction (e.g., ingress or egress). As such, a line card having both ingress and egress connectivity will have both an ingress channel and an egress channel. Thus, for

example, the ingress channel 301_1 and the egress channel 312_1 may coexist upon the same line card.

[0017] In other distributed switch architecture embodiments, each ingress channel 301_1 through 301_x transmits less than all of its ingress traffic to each egress channel 312_1 through 312_x. For example, referring again to Figure 3, if ingress channel 301_1 receives n STS-1 signals from its corresponding network line 303_1 - less than all n of the STS-1 signals received by the ingress channel 301_1 are transmitted across the node's backplane 305 over each of its output lines 306, 310, 311, ... 312.

[0018] This reduced amount of backplane 305 traffic (as compared to the "full mesh" approach described above) may be achieved by designing some degree of switching intelligence into the ingress channels 301_1 through 301_x themselves. As a result, each ingress channel 301_1 through 301_x transmits a subset of all n of the STS-1 signals it receives. Note that the actual electrical signal lines used to transport STS-1 signals across the backplane 305 may also vary from embodiment (regardless of the backplane is a full mesh backplane or is less than a full mesh backplane). For example, in one approach STS-1 signals from the same frame may be transported in parallel (e.g., as suggested by Figure 3) or in series (e.g., STS-1 signals from the same frame are transported in a time division multiplexed fashion).

BRIEF DESCRIPTION OF THE DRAWINGS

[0019] The present invention is illustrated by way of example, and not limitation, in the Figures of the accompanying drawings in which.

[0020] **Figure 1** shows a point to point SONET networking perspective.

[0021] **Figure 2** shows a ring SONET networking perspective.

[0022] **Figure 3** shows a full mesh distributed switch architecture.

[0023] **Figure 4** shows an embodiment of an egress channel.

[0024] **Figure 5** shows a configuration methodology for a node that utilizes the egress channel design of Figure 4.

[0025] **Figure 6** shows a trio of nodes that each utilize the egress channel design of Figure 4.

[0026] **Figure 7** shows a methodology for implementing protection within a SONET network.

[0027] **Figure 8** shows an embodiment of an egress channel having path termination.

DETAILED DESCRIPTION

[0028] A channel is described that has a backplane interface unit that selects a signal from a backplane. The backplane interface unit is coupled to a cross connect table that provides an indication where the signal may be found on the backplane. The indication is correlated to a logical label. The logical label is correlated to a frame location that the selected signal is transmitted within. Other appratti as well as related methods are described in more detail below.

Egress Channel Architecture

[0029] A SONET node architecture that supports a wide selection of different redundancy schemes provides customers with a wide variety of protection options. Figure 4 shows an embodiment 412 of an egress channel that at least supports 1+1, 1:N, UPSR and BLSR based redundancy schemes. The egress channel 412 of Figure 4 may be used within the full mesh distributed switch architecture, or less than a full mesh architecture, as described with respect to Figure 3.

[0030] Note that STS signals are electrical while OC signals are optical. For the sake of simplicity, the term STS is used throughout the following description although the teachings relate to both OC signals and STS signals. Furthermore, outside the United States, STM signals are used within SDH frames (rather than STS signals within SONET frames). Again, for the sake of simplicity, the term STS is used throughout the following description even though the teachings apply to either STS signals or STM signals (as well as SDH frames or SONET frames).

[0031] In the egress channel embodiment 412 of Figure 4, the backplane interface unit 401 receives a plurality of STS-1 signals from the node's backplane. For example, if the egress channel embodiment 412 of Figure 4 is implemented as the egress channel 312 of Figure 3, inputs 406_1 through 406_x of Figure 4 correspond to inputs 306_1 through 306_x of Figure 3. As such, in such an application, the backplane interface unit 401 receives all xn STS-1 signals received by the node.

[0032] Of these xn STS-1 signals, the backplane interface unit 401 may select up to n STS-1 signals for transmission over the outbound network line 404. As such, up to n STS-1 signals may be presented at the output 440 of the backplane interface unit 401. The framing unit 402 places each STS-1 signal selected by the backplane interface unit into a location within a SONET frame. Alternatively, the framing unit 402 provides SDH framing for an STM-1 signal. Further still, the backplane signals may corresponds to STS-N or STM-N signals where N is greater than 1. The line interface unit 403 transmits the framed, selected STS-1 signals over the outbound network line 404.

[0033] The line interface unit 403 may include an optic or electrical transmitter depending on whether the outbound network line 404 is an optical or electrical cable. Note that the outbound network line 404 may be implemented with a plurality of outbound network lines in a "striped" fashion. For example, an OC-48 outbound network line 404 may actually be implemented with four OC-12 outbound network lines. Thus, in various embodiments, the line interface unit

403 and the framing unit 402 may be designed to distribute the traffic from the backplane over a plurality of outbound lines.

[0034] Data provided on the selection input 450 determines which of the xn STS-1 signals are selected by the backplane interface unit 401 for transmission over outbound network line 404. In the embodiment of Figure 4, physical identification (PID) values are presented to the backplane interface unit 401 so that the appropriate STS-1 signals are selected. For example a unique PID value may be configured for each STS-1 signal presented to the backplane interface unit. Each PID value indicates where its corresponding STS-1 signal may be found on the backplane.

[0035] For example, the backplane may be designed to correlate each ingress STS-1 signal with the particular ingress channel and SONET frame location it arrived on. With such a design point, a particular backplane location or connection (e.g., a pin that an egress line card mates to) may be identified with a particular ingress STS-1 signal. For example, a first STS-1 signal may be directed to a first backplane pin that an egress line card mates to, a second STS-1 signal may be directed to a second backplane pin that an egress line card mates to, etc.

[0036] In one embodiment, the selection input 450 corresponds to a data bus that can hold up to n PID values; and, the backplane interface unit 401 is configured to provide the corresponding STS-1 signal at its output 440 for each PID value it receives. As such, the appropriate STS-1 signals are presented to

the framing unit 401 by feeding the selection input 450 with the appropriate PID values.

[0037] In an embodiment, each PID value has two components that reflect the source of the STS-1 signal to which the PID value refers. For example, a first component may reflect which backplane card slot the ingress channel that received the STS-1 signal is "plugged into"; and, a second component may reflect which location within the incoming SONET frame structure (e.g., as it exists on the incoming network line such as incoming network line 303_1 of Figure 3) that the STS-1 signal is being received within.

[0038] As an example of such a PID strategy, if a node is designed to handle a maximum of twenty four line cards that each can provide a maximum of forty eight ingress STS-1 signals - then the node may be configured to have 1152 different PID values. Each PID value may be configured to have a first component that is five bits wide (so that the full range of twenty four line card slots can be uniquely expressed) and a second component that is six bits wide (so that the full range of forty eight ingress STS-1 signals can be uniquely expressed).

[0039] As seen in the egress channel embodiment of Figure 4, the PID values are listed within a cross connect table 411 that is associated with a selection unit 405. The selection unit 405 is responsible for providing the appropriate PID values along the selection input 450 in response to configuration commands provided by the node's APS function 422 and the node's Configuration function 423. The APS function 422, as described in the background, controls the "switch

over" from working resources to protection resources in the event of a failure or degradation event.

[0040] The Configuration function 423 establishes the switching fabric of the node (e.g., establishes which egress STS-1 signal a particular ingress STS-1 signal is converted into). Both the APS 422 and Configuration 423 functions are typically implemented in software and are shown as different functions in Figure 4 (however, in other embodiments, they may be combined together and viewed as a single function). The APS 422 and Configuration functions 423 may be distributed (e.g., executed upon a plurality of line cards) or centralized (e.g., executed upon a "motherboard" that plugs into the node's backplane and is communicatively coupled to each line card in the node so that configuration/APS commands may be directed to them).

[0041] A more detailed explanation of an embodiment of how the APS and Configuration functions 422, through manipulation of the selection unit 405, may control which STS-1 signals are selected from the backplane interface unit 401 is described immediately below. The egress channel embodiment 412 of Figure 4 is designed according to the following perspective: 1) a unique physical ID (PID), as described above, is established for each ingress STS-1 signal of the node; and 2) a unique logical ID (LID) is established for each egress STS-1 signal of the node.

[0042] By correlating a PID to an LID, a component of the node's switching fabric is established. That is, the correlation of a particular PID to a particular LID corresponds to a description of the conversion of a particular ingress STS-1

signal to a particular egress STS-1 signal. A full description of the node's switching fabric may be created by establishing a PID/LID value for each ingress STS-1 signal received by the node.

[0043] In the egress channel embodiment 412 of Figure 4, the cross connect (CC) table 411 of Figure 4 corresponds to such a description. That is, CC table 411 of Figure 4 provides a complete listing of the PID/LID correlation for each STS-1 signal that is switched by the node. In an embodiment, the CC table 411 is implemented as a "look-up" table that provides an output PID value (along the selection input 450) for each input LID presented to the table 411.

[0044] As already described above, a plurality of output PID values are provided to the backplane interface unit 401 so that the appropriate STS-1 signals are selected for transmission over outbound transmission line 404. As such, a plurality of corresponding LID values are provided to the CC table 411 as a look-up input parameter.

[0045] Recall that a unique LID value is established for each of the node's egress STS-1 signals. As the egress channel 412 embodiment of Figure 4 can handle the transmission of n such egress STS-1 signals, egress table 410 corresponds to a correlation of the LID value provided to each of the egress STS-1 signals transmitted by the egress channel 412. That is, in the exemplary embodiment of Figure 4, the framing unit 402 can provide framing for a maximum number of "n" egress STS-1 signals.

[0046] As such, if the full capacity of the framing unit 402 is to be utilized, the right hand side of the egress table 410 corresponds to a listing (e.g., STS_1, STS_2,

STS_3, ...STS_n) of each framing location that may be crafted by the framing unit 402. As each of these framing locations corresponds to an egress STS-1 signal, each are provided with a unique LID value (e.g., LID_{12}, LID_{24}, LID_{31}, ... LID_{115}). Furthermore, as the CC table 411 corresponds to a full description of the node's switch fabric, the presentation to the CC table 411 of each LID value within the egress table 410 (as an input parameter for a look up) will provide the proper selection of STS-1 signals for framing by the framing unit 402.

[0047] During a configuration mode, the Configuration function 423 configures the contents of the CC table 411 and the egress table 410. By configuring an egress table within each egress channel within a node, the Configuration function 423 effectively describes which STS-1egress signals are to be transmitted from which egress line card.

[0048] Furthermore, by configuring the CC table 410, the Configuration function 423 effectively describes the conversion of each ingress STS-1 signal into its corresponding egress STS-1 signal (which effectively describes the node's switching fabric). As the Configuration function 423 is able to configure both the CC table 410 and the egress table 411, note that Figure 4 indicates that the Configuration function is coupled to both of these tables.

[0049] Figure 5 shows an embodiment of a configuration methodology 500 that corresponds to the description provided just above. That is, as described above, a logical identification label (LID) is assigned 501 to each of a node's egress STS-1 signals. Each of the logical identification labels are also correlated 502 to a particular outbound SONET frame and SONET frame location. An

embodiment of processes 501 and 502 includes configuring the egress table 410 of Figure for each egress channel within the node. This activity corresponds to defining the outbound networking line and the SONET frame location for each egress STS-1 signal that is transmitted by the node.

[0050] Once the node's egress STS-1 signals have been defined as described above, a physical identification label (PID) is assigned 503 to each of a node's ingress STS-1 signals. The node's switching fabric is then configured 504 by correlating each physical identification label to a logical identification label. That is, the conversion of each ingress STS-1 signal into an egress STS-1 signal has been mapped. An embodiment of processes 503 and 504 includes configuring the CC table 411 of Figure 4.

Protection Implementation

[0051] The egress channel architecture 412 of Figure 4 also allows for easy implementation of the various protection schemes that are discussed in the background. Figure 6 shows an exemplary implementation of a trio of nodes 631, 632 and 633. Nodes 631 and 632 are communicatively coupled by networking lines 604_1 through 604_6. Nodes 632 and 633 are communicatively coupled by networking lines 604_1 through 604_6. Notice that the nodes 631, 632 and 633 of Figure 6 utilize the egress channel architecture 412 of Figure 4.

[0052] That is, referring to Figure 4 and Figure 6: 1) each of units 601_x of Figure 6 (i.e., units 601_1 through 601_{12}) correspond to the backplane interface unit 401 of Figure 4; 2) each of units 602_x of Figure 6 (i.e., units 602_1 through 602_{12}) correspond to the framing unit 402 of Figure 4; 3) each of units 603_x of

Figure 6 (i.e., units 603_1 through 603_{12}) correspond to the line interface unit 403 of Figure 4; 4) each of units 611_x of Figure 6 (i.e., units 611_1 through 611_{12}) correspond to the cross connect table 411 of Figure 4; and 5) each of units 610_x of Figure 6 (i.e., units 610_1 through 610_{12}) correspond to the egress table 410 of Figure 4.

[0053] Within Figure 6: 1) line cards 612_1 through 612_3 are "plugged into" the backplane of node 631; 2) line cards 612_4 through 612_9 are "plugged into" the backplane 605 of node 632; and 3) line cards 612_{10} through 612_{12} are "plugged into" the backplane of node 633. The backplane structure of nodes 631 and 633 (as well as other line cards that may be plugged into these backplane structures) are not shown in Figure 6 for convenience.

[0054] Note that there is a distinction between a line card and a channel. A line card is a card that can be coupled to a network line. A channel is a data path within a line card that handles traffic flow in a particular direction (e.g., ingress or egress). As such, a line card having both ingress and egress connectivity (such as line cards 612_1 through 612_{12} of Figure 6) will have both an ingress channel and an egress channel. Thus referring briefly back to Figure 3 and recalling its discussion, for example, the ingress channel 301_1 and egress channel 312_1 may coexist upon the same line card.

[0055] The ingress and egress channels of a single line card may be drawn to be overlapping as seen in Figure 6. Specifically, for convenience, the framing units 602_1 through 602_{12} of Figure 6 and the line interface units 603_1 through 603_{12} of Figure 6 are drawn as having both ingress and egress functionality.

Note that, similar to the discussion provided above with respect to Figure 4 and the striping of a plurality of network lines, the line interface units 603_1 through 603_{12} and framing units 602_1 through 602_{12} may be designed to handle a plurality of ingress lines.

[0056] As such, as seen in Figure 6, each framing unit 602_1 through 602_{12} and each line interface unit 603_1 through 603_{12} has an ingress input and output as well as an egress input and output. The exemplary implementation of Figure 6 will be used in the discussion that follows to demonstrate examples of 1+1, 1:N, BLSR and UPSR redundancy schemes that are implemented with the egress channel architecture 412 of Figure 4.

[0057] As an example of 1+1 protection, assume network line 604_1 and network line 604_3 correspond to a 1+1 protection scheme that is established for traffic flowing to the right (as seen in Figure 6) between nodes 631 and 632. Specifically, assume that network line 604_1 corresponds to the working network line and that network line 604_3 corresponds to the protection network line.

[0058] In an embodiment where node 632 acts as a simple repeater, the 1+1 protection scheme between nodes 631 and 632 may be repeated between nodes 632 and 633. That is, the working traffic received by node 632 on working network line 604_1 is retransmitted as working traffic on networking line 604_7. Similarly, the protection traffic received by node 632 on protection network line 604_3 is retransmitted as protection traffic on networking line 604_{10}.

[0059] In order to implement this configuration, the Configuration function 623_2 of node 632 configures the CC table 611_7 and egress table 610_7 of line card 612_7

such that its corresponding backplane interface unit 601_7 selects each ingress STS-1 signal received by line card 612_4. Similarly, the Configuration function 623_2 of node 632 also configures the CC table 611_{10} and egress table 610_{10} of line card 612_{10} such that its corresponding backplane interface unit 601_{10} selects each ingress STS-1 signal received by line card 612_5.

[0060] Note that all the CC tables within a node (e.g., CC tables 611_1 through 611_3 within node 631 and CC tables 611_4 through 611_9 within node 632) may be identically maintained through a node. Thus, a change to a CC table within a particular line card (such as the change described just above with respect to CC table 611_7 and line card 612_7) may be accompanied by an identical change to all the other CC tables within the node. As such, each line card maintains an updated “copy” of the node’s switching fabric.

[0061] Continuing then, if the working network line 604_1 between nodes 631 and 632 suffers a failure or degradation event, the APS function 622_2 of node 632 changes the CC table 611_7 of line card 612_7 such that its corresponding backplane interface unit 601_7 selects each STS-1 signal received by line card 612_5 rather than line card 612_4. As the LID values provided by the egress table 610_7 remain unchanged, the node 632 will “switch” from working traffic to protection traffic by replacing (within the CC table 611_7) the original PID values (that corresponded to the ingress STS-1 signals received by line card 612_4) with PID values that correspond to the ingress STS-1 signals received by line card 612_5,.

[0062] The APS function 622_2 of node 632 may then engage in communication with the APS function 622_1 of node 631 (e.g., via the exchange of K bytes) to establish another active line between the two nodes 631, 632. For example, the pair of APS functions 622_1, 622_2 may agree that network line 604_5 should be utilized as the protection line for new working line 604_3.

[0063] As such, APS function 622_2 configures the CC table 611_3 of line card 612_3 (e.g., by replacing PID values as described above) such that its backplane interface unit 601_3 selects the ingress STS-1 signals (relative to node 631) that were originally selected by the backplane interface unit 601_1 of line card 612_1 before the failure/degradation of line 604_1. In a similar manner, the APS function 622_2 of node 632 may reconfigure the CC table 611_8 of line card 612_8 so that its backplane interface unit 601_8 selects the STS-1 signals now being received at line card 612_6.

[0064] As an example of 1:N protection where node 632 acts as a repeater, assume network lines 604_1, 604_2 and 604_3 correspond to a 1:N protection scheme that is established for traffic flowing to the right (as seen in Figure 6) between nodes 631 and 632. Specifically, assume that network lines 604_1 and 604_2 correspond to working network lines and that network line 604_3 corresponds to the protection network line.

[0065] If either of the working lines 604_1 and 604_2 suffer a failure or degradation event, the APS function 622_1 of node 631 reconfigures the CC table 611_3 of line card 612_3 so that the STS-1 signals originally carried by the failed/degraded working line begin to be carried on the protection line 604_5.

Similarly, the APS function 622_2 of node 632 reconfigures its CC tables so that the ingress STS-1 signals originally received on the failed/degraded working line begin to be selected from line card 612_6 (in order to implement their conversion into an egress STS-1 signal from node 632) rather than line card 612_4 or line card 612_5.

[0066] As an example of UPSR or BLSR protection where node 632 acts as a repeater, network lines 604_1, 604_4, 604_8 and 604_9 may correspond to links in a ring protection scheme for nodes 631, 632 and 633. Specifically, as an example, network lines 604_1 and 604_8 may correspond to a working direction while network lines 604_9 and 604_4 correspond to a protection direction. As such, the CC tables of node 632 are set such that line card 612_7 re-transmits the STS-1 signals received on line card 612_4; and, line card 612_5 re-transmits the STS-1 signals received on line card 612_8.

[0067] Recall from the background that a form of path protection includes the ability to provide protection for STS-1 signals on an individual basis rather than de facto protection for all the STS-1 signals that exist on a particular line. The egress channel architecture 412, 612 observed in Figures 4 and 6 allows for path protection as described above. For example, a "path" level protection line card may be created by configuring its corresponding CC table to select certain STS-1 signals even though their corresponding working STS-1 signals are transmitted on different outbound transmission lines. That is, a first protection STS-1 signal is selected and a second protection STS-1 signal is selected where their

corresponding working STS-1 signals are transmitted from the node as working traffic on different outbound network lines.

[0068] Figure 7 shows a methodology 700 that encompasses various embodiments (such as those discussed just above with respect to Figure 6) that relate to "switching over" from a working resource to a protection resource in light of a failure or degradation event. Note that the initial sequences 701, 702, 703, 704 of the methodology 700 of Figure 7 correspond to the configuration methodology 500 originally presented in Figure 5.

[0069] As such, following the initial configuration of a node, the node operates 705 (i.e., actively switches or routes ingress signals into egress signals). If a failure or degradation event 706 does not occur, the node continues to operate 705. However, upon the detection of a line failure or degradation event 706, a replacement physical identification (PID) label is found 707 for each affected PID. An affected PID corresponds to a PID that describes an ingress signal who has suffered the failure or degradation event.

[0070] For example, if working line 604_1 of Figure 6 suffers a failure, each of the ingress signals being received at line card 612_4 are affected. As a result, the PID values that correspond to these affected ingress signals may be referred to as affected PID values. Recall from the various recovery examples provided above with respect to Figure 6 that protection traffic effectively replaced working traffic by replacing the affected PID values (within a CC table) with new PID values that corresponded to the source of the protection traffic. Such an activity

therefore corresponds to an embodiment of the activity represented by sequences 707 and 704 in Figure 7.

[0071] Note that the logical identification (LID) values may be viewed as permanent or at least quasi-permanent in the sense that a failure or degradation event may be recovered from without implementing a change to any LID value. This is consistent with the fact that each LID value corresponds to a label for each egress signal that the node has been configured to provide resources for. A failure or degradation event does not change the fact that the node has been configured to provide a certain number of egress signals. As such, changes in LID values are typically unwarranted in light of such an event.

[0072] It is important to point out that other embodiments relating to the design approach discussed herein are possible. For example, the line card design approach of Figure 6 may be extended to include Internet Protocol (IP) routing functionality. Figure 8 shows an example. In the embodiment 812 of Figure 8, the backplane selection unit 801 provides a first output 840 to a framing unit 802 as well as a second output 842 to a path termination unit 843.

[0073] In an embodiment, the path termination unit 843 includes cell assembly/reassembly and/or other path termination related functions so that an STS-1 signal interface is maintained at the second output 842 and the path termination unit output 844. In a further embodiment, the path termination unit 843 may be configured to include Internet Protocol (IP) routing or switching. Typically, STS-1 signals leaving the path termination unit 843 are selected by another line card from the backplane for conversion into an egress signal.

[0074] Path termination and/or IP switching/routing may add another dimension to the logical identification (LID) strategy. For example, the egress table 810 and CC table 811 may be expanded to include unique LID values for those STS-1 signals destined for the framing unit 802 as well as those STS-1 signals destined for the path termination unit 843. The LID values may also be configured with a header or other sub-component that identifies whether the referenced signal should be provided at the first output 840 or the second output 842.

[0075] Referring to the embodiments 412, 812 of Figures 4 and 8, note that a table (such as tables 410, 411 and 610, 611) may be implemented with a memory, register, or other data storage device. Furthermore, tables 410, 411 and 610, 611 may be viewed as any form of correlation between the enlisted entries. For example, egress table 410, 610 does not need to be implemented as a "look-up" table. Egress table 410, 610 may be embodied in any way that establishes correlation between a particular outbound frame location and a label (e.g., the LID) that identifies an egress signal.

[0076] For example, in an embodiment, the framing unit 402, 802 has a unique input for each STS-1 signal that is to be framed by the framing unit 402, 802. If LID values are provided to the CC table (e.g., where the LID value is obtained by means other than a look-up operation) that control which backplane signal is offered to which framer unit input, the sought for correlation has been established.